On 2DOF agreement protocols

Gal Barkai   Leonid Mirkin   Daniel Zelazo

Technion - Israel Institute of Technology

IAAC^3 - IAAC 2025 Control Conference

April 2025

Networked multi-agent systems

- Multiple dynamic units interacting over a network to achieve a collective goal.

Networked multi-agent systems

- Multiple dynamic units interacting over a network to achieve a collective goal.
- Many applications
  - sensor networks
  - electrical microgrids
  - multi-robot coordination

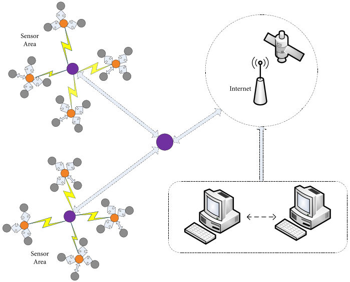

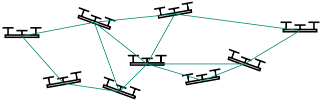

Networked multi-agent systems

- Multiple dynamic units interacting over a network to achieve a collective goal.
- Many applications
  - sensor networks
  - electrical microgrids
  - multi-robot coordination
- Often control/computation is cheap while communication is expensive.

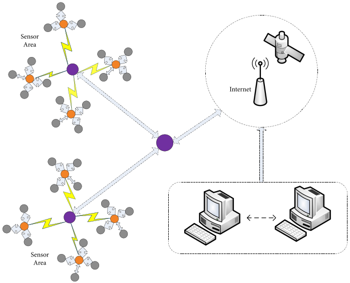

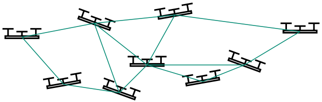

Archetype MAS control problem: agreement of integrators

An ensemble of ν independent integrator agents

$$P_i(s) = \frac{1}{s} \implies P(s) = \frac{1}{s} I_\nu$$

Goal (asymptotic agreement):

$$\lim_{t \to \infty} (y_i(t) - y_j(t)) = 0, \quad \forall i, j$$

The challenge: communication is subject to spatial constraints.

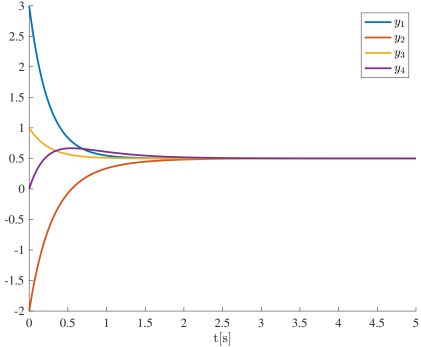

Figure: Consensus Trajectories $k \equiv 1$)

Archetype MAS control problem: agreement of integrators

An ensemble of ν independent integrator agents

$$P_i(s) = \frac{1}{s} \implies P(s) = \frac{1}{s} I_\nu$$

Goal (asymptotic agreement):

$$\lim_{t\to\infty} (y_i(t) - y_j(t)) = 0, \quad \forall i, j$$

The challenge: communication is subject to spatial constraints.

The solution: the consensus protocol

$$u_i(t) = -k \sum_{j\in\mathcal{N}_i} (y_i - y_j).$$

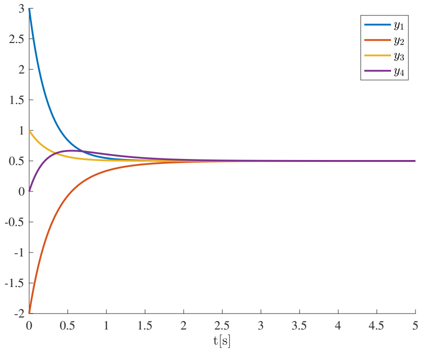

Figure: Consensus Trajectories $k \equiv 1$)

Archetype MAS control problem: agreement of integrators

An ensemble of ν independent integrator agents

$$P_i(s) = \frac{1}{s} \implies P(s) = \frac{1}{s} I_\nu$$

Goal (asymptotic agreement):

$$\lim_{t\to\infty} (y_i(t) - y_j(t)) = 0, \quad \forall i, j$$

The challenge: communication is subject to spatial constraints.

The solution: the consensus protocol

$$u_i(t) = -k \sum_{j \in \mathcal{N}_i} (y_i - y_j).$$

$$\mathcal{N}_1 = \{P_2, P_3, P_4\},\ \mathcal{N}_2 = \{P_1, P_3\}$$

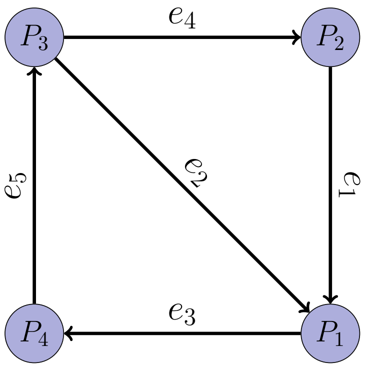

Figure: An example coupling graph

Archetype MAS control problem: agreement of integrators

An ensemble of ν independent integrator agents

$$P_i(s) = \frac{1}{s} \implies P(s) = \frac{1}{s} I_\nu$$

Goal (asymptotic agreement):

$$\lim_{t \to \infty} (y_i(t) - y_j(t)) = 0, \quad \forall i, j$$

The challenge: communication is subject to spatial constraints.

The solution: the consensus protocol

$$u_i(t) = -k \sum_{j \in \mathcal{N}_i} (y_i - y_j).$$

- Agreement is reached for all initial conditions if $\mathcal{G}$ is connected and $k > 0$.

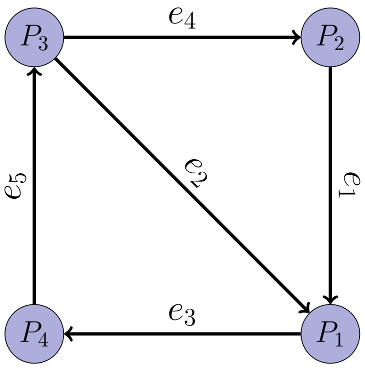

Figure: An example coupling graph

Nominal behavior of the consensus protocol

- Aggregating the protocol

$$u_i(t) = -k \sum_{j \in \mathcal{N}_i} (y_i(t) - y_j(t)) \implies u(t) = -k L_{\mathcal{G}} y(t), \quad L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}$$

  where $L_{\mathcal{G}}$, $D_{\mathcal{G}}$, and $A_{\mathcal{G}}$ are called the graph Laplacian, Degree, and Adjacency matrices.

- The Laplacian, and hence $\mathcal{G}$, determines the dynamics.

Nominal behavior of the consensus protocol

- Aggregating the protocol

$$u_i(t) = -k \sum_{j \in \mathcal{N}_i} (y_i(t) - y_j(t)) \implies u(t) = -kL_{\mathcal{G}}y(t), \quad L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}$$

  where $L_{\mathcal{G}}$, $D_{\mathcal{G}}$, and $A_{\mathcal{G}}$ are called the graph Laplacian, Degree, and Adjacency matrices.

- The Laplacian, and hence $\mathcal{G}$, determines the dynamics.
  - $L_{\mathcal{G}}$ is PSD and $L_{\mathcal{G}}\mathbb{1} = 0$.
  - If $\mathcal{G}$ is connected, 0 is a simple eigenvalue of $L_{\mathcal{G}}$.
  - The projection on the agreement manifold, y_{agr}, is constant and independent of k.
  - Convergence rate is exponential in k.

Nominal behavior of the consensus protocol

- Aggregating the protocol

$$u_i(t) = -k \sum_{j \in \mathcal{N}_i} (y_i(t) - y_j(t)) \implies u(t) = -kL_{\mathcal{G}} y(t), \quad L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}$$

  where $L_{\mathcal{G}}$, $D_{\mathcal{G}}$, and $A_{\mathcal{G}}$ are called the graph Laplacian, Degree, and Adjacency matrices.

- The Laplacian, and hence $\mathcal{G}$, determines the dynamics.
  - $L_{\mathcal{G}}$ is PSD and $L_{\mathcal{G}} \mathbb{1} = 0$.
  - If $\mathcal{G}$ is connected, 0 is a simple eigenvalue of $L_{\mathcal{G}}$.
  - The projection on the agreement manifold, y_{agr}, is constant and independent of k.
  - Convergence rate is exponential in k.
- This is not limited to integrators.

General consensus protocols

Consider ν identical SISO agents with dynamics P

$$u_i = -R_0 \sum_{j \in \mathcal{N}_i} (y_i - y_j) \quad \Longrightarrow \quad y = \left(I_\nu - P R_0 L_{\mathcal{G}}\right)^{-1} P y_0$$

where y_0 are the initial conditions and R_0 a possibly dynamic controller.

General consensus protocols

Consider ν identical SISO agents with dynamics P

$$u_i = -R_0 \sum_{j \in \mathcal{N}_i} (y_i - y_j) \quad \Longrightarrow \quad y = \left(I_\nu - PR_0 L_{\mathcal{G}}\right)^{-1} P y_0$$

where y_0 are the initial conditions and R_0 a possibly dynamic controller.

- The Laplacian is diagonalizable

$$Ty = \tilde{y} = \mathrm{diag}\{(1 - \lambda_i P R_0)^{-1} P\}\tilde{y}_0, \quad T L_{\mathcal{G}} T^{-1} = \mathrm{diag}\{\lambda_i\},$$

  and the transformed agents are decoupled.

General consensus protocols

Consider ν identical SISO agents with dynamics P

$$u_i = -R_0 \sum_{j \in \mathcal{N}_i} (y_i - y_j) \quad \Longrightarrow \quad y = \left(I_\nu - PR_0 L_\mathcal{G}\right)^{-1} P y_0$$

where y_0 are the initial conditions and R_0 a possibly dynamic controller.

- The Laplacian is diagonalizable

$$Ty = \tilde{y} = \mathrm{diag}\{(1 - \lambda_i P R_0)^{-1} P\} \tilde{y}_0, \quad T L_\mathcal{G} T^{-1} = \mathrm{diag}\{\lambda_i\},$$

  and the transformed agents are decoupled.
- $\tilde{y}_1$ corresponds to the direction of $\mathbb{1}$ and $\lambda_1 = 0$.

General consensus protocols

Consider ν identical SISO agents with dynamics P

$$u_i = -R_0 \sum_{j \in \mathcal{N}_i} (y_i - y_j) \quad \Longrightarrow \quad y = \left(I_\nu - PR_0 L_{\mathcal{G}}\right)^{-1} P y_0$$

where y_0 are the initial conditions and R_0 a possibly dynamic controller.

- The Laplacian is diagonalizable

$$Ty = \tilde{y} = \operatorname{diag}\{(1 - \lambda_i PR_0)^{-1} P\}\tilde{y}_0, \quad TL_{\mathcal{G}}T^{-1} = \operatorname{diag}\{\lambda_i\},$$

  and the transformed agents are decoupled.
- $\tilde{y}_1$ corresponds to the direction of $\mathbb{1}$ and $\lambda_1 = 0$.

In general

1. Systems $2, \ldots, \nu$ must be stabilized.
2. Since $\lambda_1 = 0$, the agreement trajectory is set by P.

Real life isn't perfect

- Replace y_j in the consensus protocol with $y_j + n_{ij}$ for some network-induced white noise.

Real life isn't perfect

- Replace y_j in the consensus protocol with $y_j + n_{ij}$ for some network-induced white noise.

- Even for simple integrators:
  - Agreement is not reached.
  - y_{agr} is no longer constant.

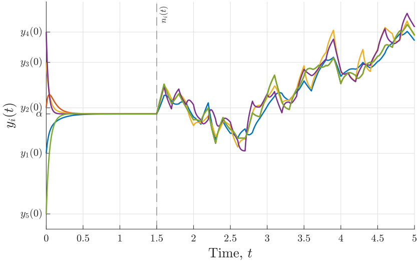

Figure: Outputs with white network noise

Real life isn't perfect

- Replace y_j in the consensus protocol with $y_j + n_{ij}$ for some network-induced white noise.
- Even for simple integrators:
  - Agreement is not reached.
  - y_{agr} is no longer constant.
- Inherent tradeoff
  - Larger $k \implies$ better nominal convergence.
  - Larger $k \implies$ larger noise sensitivity.

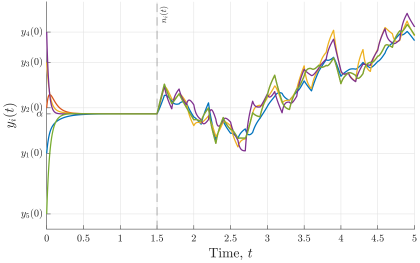

Figure: Outputs with white network noise

Real life isn't perfect

- Replace y_j in the consensus protocol with $y_j + n_{ij}$ for some network-induced white noise.
- Even for simple integrators:
  - Agreement is not reached.
  - y_{agr} is no longer constant.
- Inherent tradeoff
  - Larger $k \implies$ better nominal convergence.
  - Larger $k \implies$ larger noise sensitivity.
- This behavior is persistent even for non integrator agents and arbitrary LTI controller R_0.

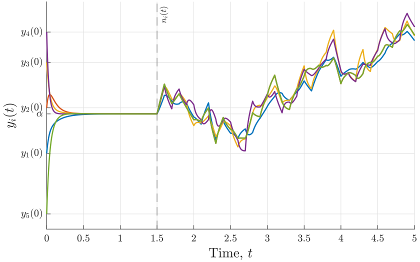

Figure: Outputs with white network noise

A feedback perspective

We can rewrite the local consensus protocol as

$$u_i(t) = -k|\mathcal{N}_i|\big(\underbrace{y_i(t) - \tilde{r}_i(t)}_{e_i}\big), \quad \tilde{r}_i(t) \coloneqq \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} y_j(t).$$

- Each agent attempts to track the *centroid* of their neighbors!

A feedback perspective

We can rewrite the local consensus protocol as

$$u_i(t) = -k|\mathcal{N}_i|\big(\underbrace{y_i(t) - \tilde{r}_i(t)}_{e_i}\big), \quad \tilde{r}_i(t) \coloneqq \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} y_j(t).$$

- Each agent attempts to track the *centroid* of their neighbors!
- Note, the noise affects only $\tilde{r}_i$, but the controller acts uniformly on the entire error.

A feedback perspective

We can rewrite the local consensus protocol as

$$u_i(t) = -k|\mathcal{N}_i|\big(\underbrace{y_i(t) - \tilde{r}_i(t)}_{e_i}\big), \quad \tilde{r}_i(t) \coloneqq \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} y_j(t).$$

- Each agent attempts to track the *centroid* of their neighbors!
- Note, the noise affects only $\tilde{r}_i$, but the controller acts uniformly on the entire error.

Classical control is the art of balancing performance and robustness, can it help us gain insight for agreement problems?

Those who forget their history are condemned to repeat it

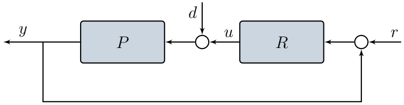

(a) UNITY-FEEDBACK SETUP

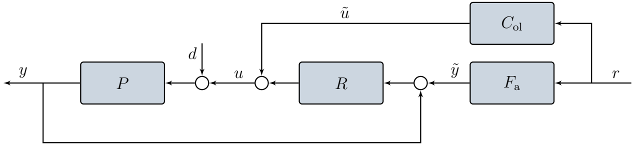

(b) 2DOF SETUP

- Unity-feedback: Let R be stabilizing, the output is given by

$$y = (I - PR)^{-1} PR\, r + (I - PR)^{-1} P\, d.$$

Those who forget their history are condemned to repeat it

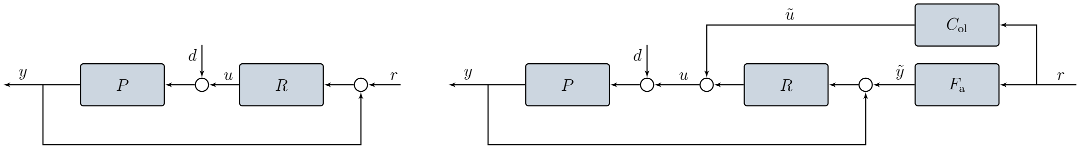

(a) Unity-feedback setup (b) 2DOF setup

- Unity-feedback: Let R be stabilizing, the output is given by

$$y = (I - PR)^{-1} PR\, r + (I - PR)^{-1} P\, d.$$

- 2DOF: Let R be stabilizing and $\tilde{y} = F_{\mathrm{a}} r = PC_{\mathrm{ol}} r$, the output is given by

$$y = F_{\mathrm{a}}\, r + (I - PR)^{-1} P\, d.$$

Those who forget their history are condemned to repeat it

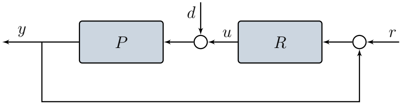

(a) UNITY-FEEDBACK SETUP

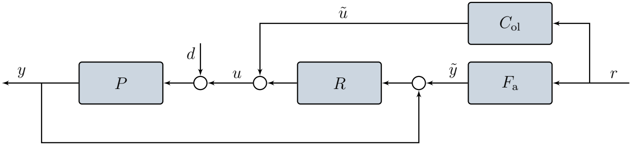

(b) 2DOF SETUP

- Unity-feedback: Let R be stabilizing, the output is given by

$$y = (I - PR)^{-1} PR\, r + (I - PR)^{-1} P\, d.$$

- 2DOF: Let R be stabilizing and $\tilde{y} = F_{\mathsf{a}} r = PC_{\mathsf{ol}} r$, the output is given by

$$y = F_{\mathsf{a}}\, r + (I - PR)^{-1} P\, d.$$

- Two-degrees-of-freedom control can decouple the sensitivity and tracking objectives!
  - This has been known since the 50's (Lang and Ham, 1955).

A 2DOF agreement protocol

- Consider the "natural reference" given by

$$\tilde{r}_i(t) := \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i}(y_j(t)) \implies \tilde{r}(t) = (A^\star_{\mathcal{G}} \otimes I)(y(t))$$

  where $A^\star_{\mathcal{G}} = D^{-1}_{\mathcal{G}} A_{\mathcal{G}}$.

- Note, $\tilde{r}_i$ is not exogenous, it is an additional network feedback.

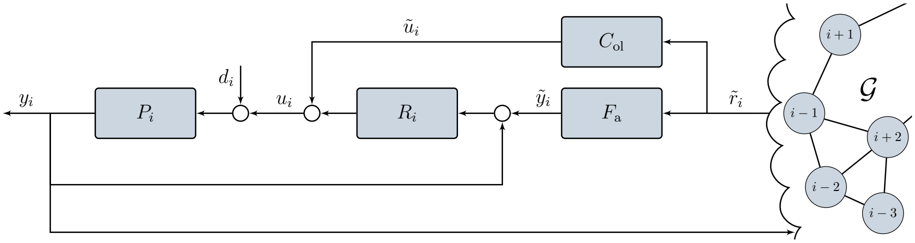

A 2DOF agreement protocol

- Consider the "natural reference" given by

$$\tilde{r}_i(t) \coloneqq \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} (y_j(t) + n_{ij}) \implies \tilde{r}(t) = (A^\star_{\mathcal{G}} \otimes I)(y(t) + n)$$

  where $A^\star_{\mathcal{G}} = D^{-1}_{\mathcal{G}} A_{\mathcal{G}}$.

- Note, $\tilde{r}_i$ is not exogenous, it is an additional network feedback.

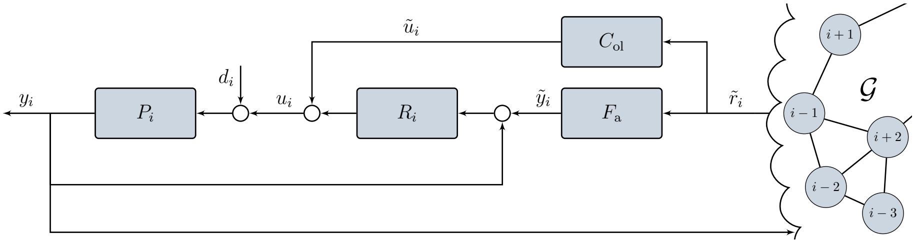

The resulting dynamics

It is easy to show that the dynamics become

$$y = (I - A^{\star}_{\mathcal{G}} \otimes F_{\mathrm{a}})^{-1} T_d d + (I - A^{\star}_{\mathcal{G}} \otimes F_{\mathrm{a}})^{-1} (A^{\star}_{\mathcal{G}} \otimes F_{\mathrm{a}}) n$$

where

- $A^{\star}_{\mathcal{G}}$ depends only on the graph (normalized adjacency matrix)
- F_{a} is an independent design parameter,
- and $T_d = \mathrm{diag}\left\{(I - P_i R_i)^{-1} P_i\right\}$ is decoupled from both the graph and F_{a}.

Two questions:

The resulting dynamics

It is easy to show that the dynamics become

$$y = (I - A_{\mathcal{G}}^{\star} \otimes F_{\mathsf{a}})^{-1} T_d d + (I - A_{\mathcal{G}}^{\star} \otimes F_{\mathsf{a}})^{-1} (A_{\mathcal{G}}^{\star} \otimes F_{\mathsf{a}}) n$$

where

- $A_{\mathcal{G}}^{\star}$ depends only on the graph (normalized adjacency matrix)
- F_{a} is an independent design parameter,
- and $T_d = \operatorname{diag}\left\{(I - P_i R_i)^{-1} P_i\right\}$ is decoupled from both the graph and F_{a}.

Two questions:

- Do the agents reach agreement?
- What is the agreement trajectory?

Agreement result

Theorem

Assume that $\mathcal{G}$ is undirected and connected, $n = 0$, and $d = y_0\delta(t)$. The agents reach asymptotic agreement if and only if

1. *each local controller R_i stabilizes its corresponding plant P_i,*
2. $\left(I_p - F_a\right)^{-1}$ *has all poles in the closed left half-plane,*
3. *and*

$$\left(I_p - \lambda_i F_a\right)^{-1} \in H_\infty, \quad \forall \lambda_i \in \operatorname{spec} A_{\mathcal{G}}^{\star} \setminus \{1\}.$$

Agreement result

Theorem

Assume that $\mathcal{G}$ is undirected and connected, $n = 0$, and $d = y_0\delta(t)$. The agents reach asymptotic agreement if and only if

1. *each local controller R_i stabilizes its corresponding plant P_i,*
2. $\left(I_p - F_a\right)^{-1}$ *has all poles in the closed left half-plane,*
3. *and*
$$\left(I_p - \lambda_i F_a\right)^{-1} \in H_\infty, \quad \forall \lambda_i \in \operatorname{spec} A_{\mathcal{G}}^{\star} \setminus \{1\}.$$

- Note that $\operatorname{spec} A_{\mathcal{G}}^{\star} \in [-1, 1]$, and $\lambda_1 = 1$ is simple and corresponds to the eigenvector $\mathbb{1}$.

Almost separation

- Agreement via 2-step design for F_{a}
  - Interpolation constraints: $(I - F_{\mathsf{a}})^{-1}$ has certain closed-loop poles
  - Robust control: $(I - \lambda_i F_{\mathsf{a}})^{-1}$ is stable $\forall \lambda_i \in [-1, 1)$.

Almost separation

- Agreement via 2-step design for F_{a}
  - Interpolation constraints: $(I - F_{\mathsf{a}})^{-1}$ has certain closed-loop poles
  - Robust control: $(I - \lambda_i F_{\mathsf{a}})^{-1}$ is stable $\forall \lambda_i \in [-1, 1)$.
- Decoupled from P_i and R_i
  - Naturally accommodates heterogeneity

Almost separation

- Agreement via 2-step design for F_a
  - Interpolation constraints: $(I - F_a)^{-1}$ has certain closed-loop poles
  - Robust control: $(I - \lambda_i F_a)^{-1}$ is stable $\forall \lambda_i \in [-1, 1)$.
- Decoupled from P_i and R_i
  - Naturally accommodates heterogeneity
- Noise response depends only on the graph and F_a.
  - Can create "off the shelf" F_a with prescribed trajectory and noise attenuation.

Almost separation

- Agreement via 2-step design for F_{a}
  - Interpolation constraints: $(I - F_{\mathsf{a}})^{-1}$ has certain closed-loop poles
  - Robust control: $(I - \lambda_i F_{\mathsf{a}})^{-1}$ is stable $\forall \lambda_i \in [-1, 1)$.
- Decoupled from P_i and R_i
  - Naturally accommodates heterogeneity
- Noise response depends only on the graph and F_{a}.
  - Can create "off the shelf" F_{a} with prescribed trajectory and noise attenuation.
- Local loop (T_{d}) can improve convergence and reject disturbances.

Many complex problems become simple.

Motivating example - redux

- Prescribed nominal performance $t_{\mathsf{s}} \approx 0.85$[s].
- Independent AWGN added to network signal.
- For 1DOF consensus $k = 6$ would amplify the noise by 6.
- The metric is drift from nominal consensus value α

$$e(t) = \|y(t) - \alpha\mathbb{1}\|$$

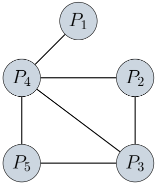

1DOF: $k = 6$

2DOF: $F_{\mathsf{a}}(s) = \frac{50}{(s+10)(s+5)}$ and $R = 50$

Metric: Distance from nominal consensus value α

$$e(t) = \|y(t) - \alpha\mathbb{1}\|$$

Motivating example - redux

- Prescribed nominal performance $t_s \approx 0.85$[s].

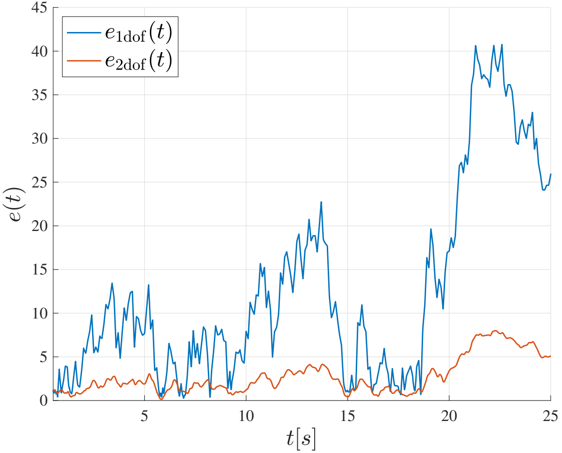

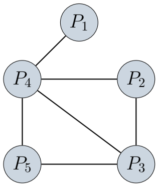

1DOF: $k = 6$

2DOF: $F_a(s) = \frac{50}{(s+10)(s+5)}$ and $R = 50$

Metric: Distance from nominal consensus value α

$$e(t) = \|y(t) - \alpha \mathbb{1}\|$$

Synchronizing integrators with $A\sin(t+\phi)$

- Impossible via standard consensus protocols
- Less obvious choice of F_a
- Local controller tuned to improve convergence rate.

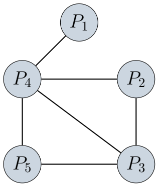

$$F_\mathsf{a}(s) = \frac{-1}{2s^4+3s^3+4s^2+3s+1}$$
$$R_0(s) = -\frac{4.142s+25}{s+10}$$

Synchronizing integrators with $A \sin(t + \phi)$

- Impossible via standard consensus protocols

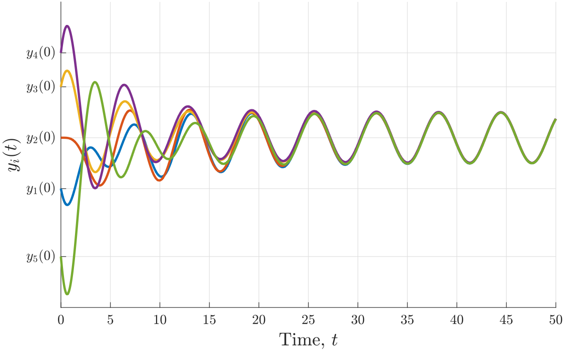

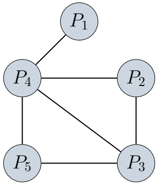

$$F_a(s) = \frac{-1}{2s^4 + 3s^3 + 4s^2 + 3s + 1}$$

$$R_0(s) = -\frac{4.142s + 25}{s + 10}$$

Consensus w/ heterogeneous agents and disturbances

- Heterogeneous agents
  $P(s) = \operatorname{diag}\left\{\frac{1}{s}, \frac{s+4}{(s+1)(s-1)}, \frac{s+2}{s(s+1)}, \frac{s+0.5}{(s+1)^2}, \frac{s+2}{s(s+1)}\right\}$
- All agents suffer from sinusodial disturbances and colored noise
- Two agents also suffer from step disturbances.

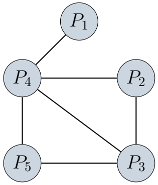

$F_a(s) = \frac{16}{(2s+1)(s^2+2\sqrt{2}s+16)}$

$R(s)$ Internal Model based.

Disturbances: Sinusoidal and step disturbances

Noise: Colored noise
$W_n(s) = 0.05s/(0.05s+1)$

Consensus w/ heterogeneous agents and disturbances

- Heterogeneous agents

$$P(s) = \mathrm{diag}\left\{\frac{1}{s}, \frac{s+4}{(s+1)(s-1)}, \frac{s+2}{s(s+1)}, \frac{s+0.5}{(s+1)^2}, \frac{s+2}{s(s+1)}\right\}$$

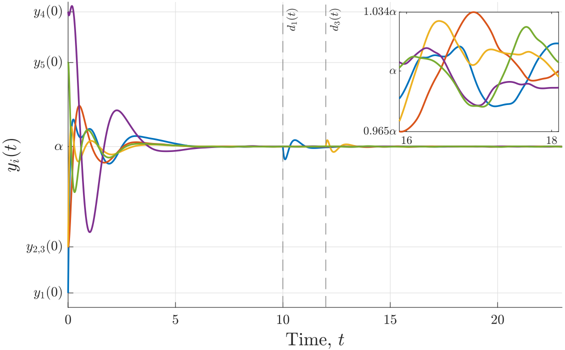

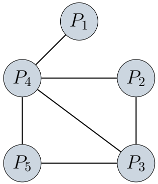

$F_a(s) = \frac{16}{(2s+1)(s^2+2\sqrt{2}s+16)}$

$R(s)$ Internal Model based.

Disturbances: Sinusoidal and step disturbances

Noise: Colored noise $W_n(s) = 0.05s/(0.05s+1)$

Concluding Remarks

- Much more in the paper
  - Design method for SISO agents
  - Robustness for graph structure
  - For static consensus: robustness to unknown heterogeneous delays, exact calculation of consensus value.
- key point: in 1DOF consensus the loop is "clopen", in 2DOF it is always closed.
- Future research:
  - Unified design method for F_{a} w/ interpolation constraints and noise attenuation.
  - Different choice of the reference $\tilde{r}$.
  - Different architectures.

Thanks for your attention!

Contact: GalBarkai@campus.technion.ac.il